Job scheduling of parametric computational mechanics studies on Cloud Computing infrastructures

Carlos García Garino Cristian Mateos Elina Pacini

HPC 2012
High Perfomance Computing, Grids and Clouds

UNCUYO UNIVERSIDAD NACIONAL DE CUYO
ITIC
UNIVERSIDAD NACIONAL DEL CENTRO DE LA PROVINCIA DE BUENOS AIRES
ISISTAN
Instituto Superior de Ingeniería de Software Tandil
CONICET-UNICEN

Outline

Motivation

- Scientists and engineers are more and more faced to the need of computational power to satisfy the ever-increasing resource intensive nature of their experiments.
- Computational Mechanics Simulation can be benefited from parallel execution of tasks in different computers in some distributed environment.

Cloud Computing

- Pleasingly parallel problems.
- Virtualization.
- Elasticity.
- Private Clouds?

Computational Mechanics

- Mature discipline. It deserves interest both in academy and industry.
- Non linear applications like finite strain elastoplastic and elasto-viscoplastic problems can be still hard to solve due to complex geometries, 3D case and strong nonlinearities.
- Finite Element Method is widely used.

Computational Mechanics Parameter Sweeping Experiments

- Input data values need to be changed.
- Material data and /or geometry data are the usual parameters considered.
- The same Finite Element Mesh is enough for most of the cases.

Computational Mechanics Parameter Sweeping Experiments

- Automatic data parameters change.
- Pleasingly parallel processing (many nonlinear FE problems).
- Automatic postprocessing.

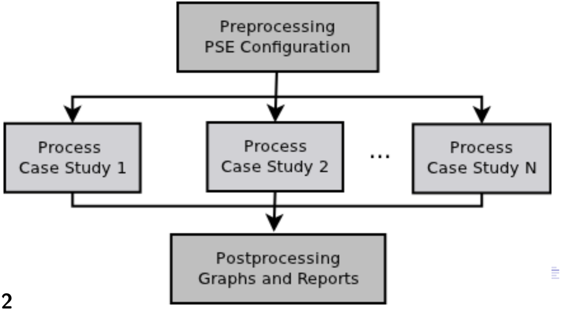

PSE Example I

Simple tensile test of a circular cylinder aluminum specimen. Changes in small imperfections.

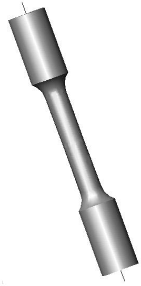

Circular Cylinder

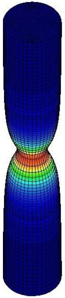

Necking Deformation

PSE Example II

Small changes in parameters value cause different answers.

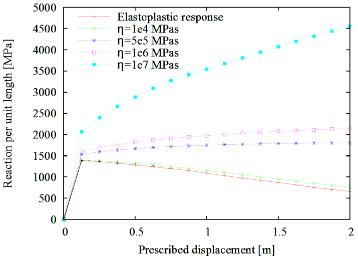

small viscosity

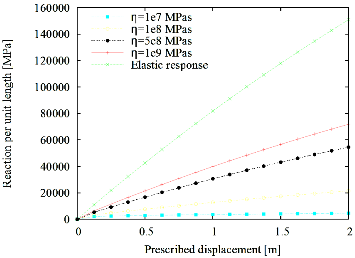

large viscosity

sensitivity of results in terms of viscosity parameter value

Deformed shapes for 2 *m* stretching

539*s*	379*s*	354*s*
$\eta = 1.e4$	$\eta = 1e6$	$\eta = 1e8$

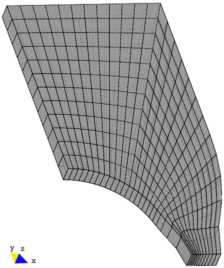

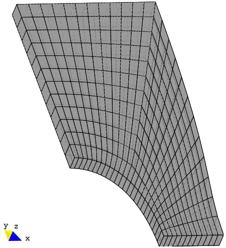

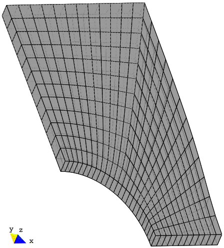

PSE domain issues

- Small changes in parameters value cause different answers.
- Different CPU times are required to complete the processes
- A proper job Scheduling becomes a requisite in order to obtain reasonable makespan.

Abstract Scheduler Architecture

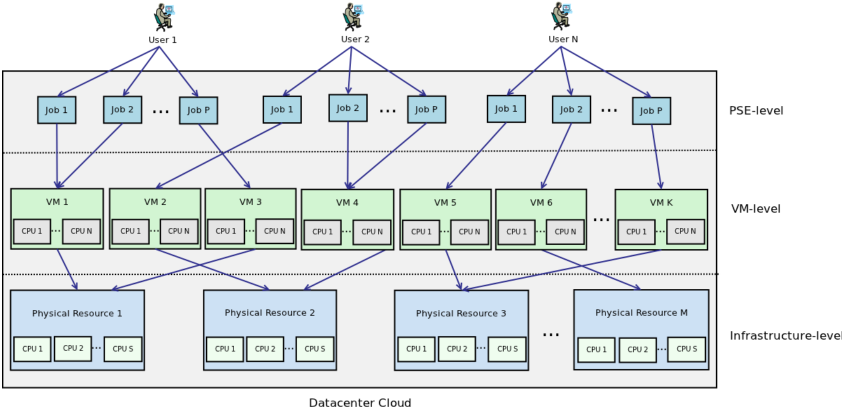

Sequence diagram of scheduling actions

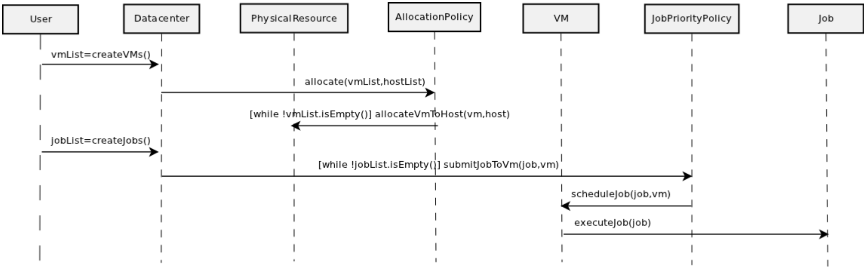

SI based Schedulers

Approaches to job scheduling based on SI: A taxonomy

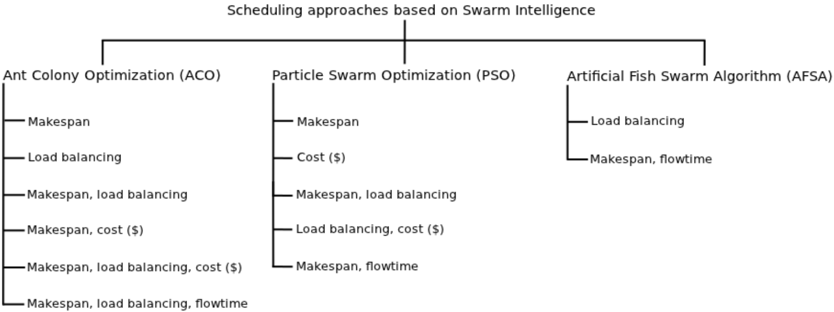

Ant Colony Optimization

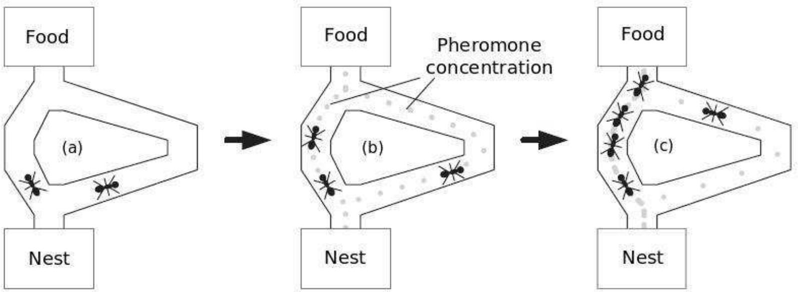

Summary

- Scheduling Policies.
  - At the host level, to assign the VMs to the physical resource, a SI scheme is used.
  - At the VM level, to assign the tasks to the VMs a priority police is considered.
- Metrics
  - 1. Flowtime $F = \sum_{j=1}^{n} (C_j(S) - A_j(S)).w_j$
  - 2. Makespan $M = max_j C_j(S)$

Case study: a viscoplastic solid PSE

The application domain under study involves a PSE of viscoplastic solids, which explore the sensitivity of solid behavior in terms of changes in the viscosity parameter η of model.

The different viscosity values of η parameter considered are: 1.10^4, 2.10^4, 3.10^4, 4.10^4, 5.10^4, 7.10^4, 1.10^5, 2.10^5, 3.10^5, 4.10^5, 5.10^5, 7.10^5, 1.10^6, 2.10^6, 3.10^6, 4.10^6, 5.10^6, 7.10^6, 1.10^7, 2.10^7, 3.10^7, 4.10^7, 5.10^7, $7.10^7$and 1.10^8 Mpa.

Finite element meshes considered

The Finite Element meshes correspond to a plain strain plate with a central circular hole and have 288 and 1152 elements. The dimensions of the plate are 18 x 10 m.

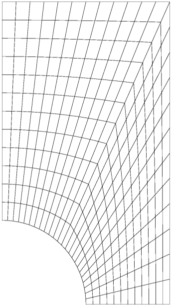

Mesh of 288 elements

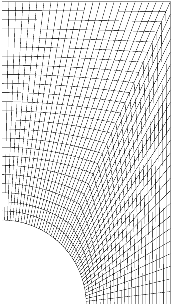

Mesh of 1152 elements

Experiment Settings

- To carry out the experiments in a real single machine we run the PSE by varying the viscosity parameter η and measuring the execution time for 25 different experiments (resulting in 25 input files with different configurations).
- The PSEs were solved using the SOGDE solver.
- The machine model is AMD Athlon(tm) 64 X2 3600+, running Ubuntu 11.04 kernel version 2.6.38-8.

CloudSim configurations

Host characteristics

Host Parameters	Value
Processing power	4,008 MIPS
Number of CPUs	2
RAM memory	4 Gbytes
Storage size	400 Gbytes
Bandwidth	100 Mbps

VM characteristics

VM Parameters	Value
Processing power	4,008 MIPS
Number of CPUs	1
RAM memory	1 Gbytes
Storage size	100 Gbytes
Bandwidth	25 Mbps
Vmm	Xen

CloudSim configuration

Parameter	Value
Number of Hosts	10
Number of VMs	40
Number of Cloudlets	from 250 to 2500

Cloudlet configuration

CloudSim-related parameters (above) and job priorities (below)

	Sogde 2D		Sogde 3D	
Cloudlet parameters	Mesh of 288	Mesh of 1152	Mesh of 288	Mesh of 1152
Length (MIPS)	52,112-104,225	244,527-469,011	216,467-268,579	1,362,938-2,160,657
PEs	1	1	1	1
Input Size (bytes)	40,038	93,082	91,136	291,738
Output Size (bytes)	722,432	2,202,010	1,677,722	5,662,310

	Sogde 2D		Sogde 3D	
Cloudlet priority	Mesh of 288	Mesh of 1152	Mesh of 288	Mesh of 1152
Low (w_i=1)	52,112-68,147	244,527-272,588	216,467-228,492	1,362,938-1,495,223
Medium (w_i=2)	80,173-96,207	280,605-348,752	232,501-236,510	1,503,240-1,607,465
High (w_i=3)	100,216-104,225	352,760-469,011	240,518-268,579	1,747,767-2,160,657

3D 1152 elements Mesh: Scalability

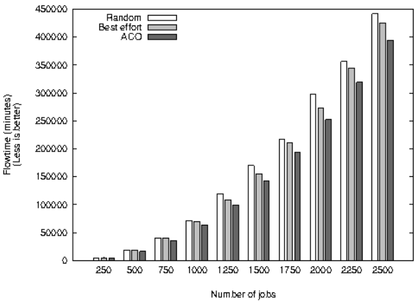

Flowtime

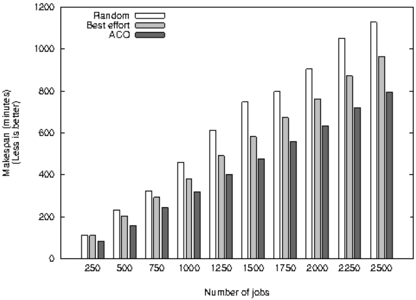

Makespan

3D 1152 elements Mesh: ACO vs Best Effort

#jobs	Mesh of 288 elements		Mesh of 1,152 elements	
	Flowtime	Makespan	Flowtime	Makespan
250	15.16	22.31	15.59	23.40
500	11.42	19.27	11.95	21.23
750	10.06	16.96	10.40	16.80
1,000	8.66	16.66	8.52	16.66
1,250	8.30	17.88	8.35	18.16
1,500	8.06	17.52	8.22	18.24
1,750	7.93	16.74	8.09	16.72
2,000	7.47	16.63	7.41	16.66
2,250	7.37	17.30	7.43	17.51
2,500	7.36	17.19	7.45	17.61

3D 1152 elements Mesh - Priority results

Scheduler	Mesh of 288 elements		Mesh of 1,152 elements	
	Flowtime (mins.)	Makespan (mins.)	Flowtime (mins.)	Makespan (mins.)
Random	107,944.01	191.95	509,950.71	1,128.37
Random (priority)	85,353.07	186.37	441,620.41	1,128.37
Gain (0-100%)	20.93	2.90	13.40	0.00
Best effort	95,510.52	148.88	491,403.19	959.88
Best effort (priority)	73,488.24	148.49	425,190.81	962.94
Gain (0-100%)	23.06	0.26	13.47	-0.32
ACO	88,729.85	123.00	456,213.40	793.37
ACO (priority)	68,079.36	122.97	393,533.76	793.32
Gain (0-100%)	23.27	0.02	13.74	0.01

3D 1152 elements Mesh - Horizontal Scalability

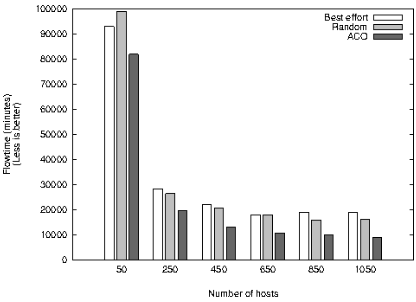

Flowtime

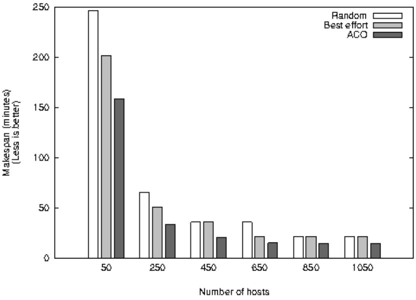

Makespan

Conclusions

- Simulated experiments were executed with the help of the CloudSim toolkit and real PSE job data.
- The proposed scheduler can effectively handle a large number of jobs.
- For Computational Mechanics PSE tested the proposed Scheduler performs well in comparison with other schemes .
- Priority considerations at Job-VM level improves flowtime.

Future Work

- Particle Swarm Optimization based scheduling algorithm is currently in progress.
- Energy consumption will be addressed. Clearly, simpler scheduling policies require fairly less resource usage, compared to more complex policies such as our algorithm. Then, flowtime/makespan vs energy consumption tradeoff will be considered.

Future Work

- Eventually, we will materialize the resulting job schedulers on top of a real (but not simulated) Cloud platform.
- A specialized Computational Mechanics Cloud Computing is a mid term goal.

Questions?

Thanks for your attention!